arXiv:1402.4312v2 [quant-ph] 16 Apr 2014

Two Results about Quantum Messages

Hartmut Klauck* Supartha Podder†

Abstract

We show two results about the relationship between quantum and classical messages. Our first contribution is to show how to replace a quantum message in a one-way communication protocol by a deterministic message, establishing that for all partial Boolean functions $f : \{0,1\}^n \times \{0,1\}^m \to \{0,1\}$ we have $D^{A\to B}(f) \leq O(Q^{A\to B,*}(f) \cdot m)$. This bound was previously known for total functions, while for partial functions this improves on results by Aaronson [Aar05, Aar07], in which either a log-factor on the right hand is present, or the left hand side is $R^{A\to B}(f)$, and in which also no entanglement is allowed.

In our second contribution we investigate the power of quantum proofs over classical proofs. We give the first example of a scenario, where quantum proofs lead to exponential savings in computing a Boolean function. The previously only known separation between the power of quantum and classical proofs is in a setting where the input is also quantum [AK07].

We exhibit a partial Boolean function f, such that there is a one-way quantum communication protocol receiving a quantum proof (i.e., a protocol of type QMA) that has cost $O(\log n)$ for f, whereas every one-way quantum protocol for f receiving a classical proof (protocol of type QCMA) requires communication $\Omega(\sqrt{n}/\log n)$.

*Division of Mathematical Sciences, Nanyang Technological University, Singapore 637371 & Centre for Quantum Technologies, National University of Singapore, Singapore 117543. E-mail: hklauck@gmail.com. This work is funded by the Singapore Ministry of Education (partly through the Academic Research Fund Tier 3 MOE2012-T3-1-009) and by the Singapore National Research Foundation.

†Centre for Quantum Technologies. E-mail: supartha@gmail.com.

1 Introduction

The power of using quantum messages over classical messages is a central topic in information and communication theory. It is always good to understand such questions well in the simplest settings where they arise. An example is the setting of one-way communication complexity, which is rich enough to lead to many interesting results, yet accessible enough for us to show results about deep questions like the relationship between different computational modes, e.g. quantum versus classical or nondeterministic versus deterministic.

1.1 One-way Communication Complexity

Perhaps the simplest question one can ask about the power of quantum messages is the relationship between quantum and classical one-way protocols. Alice sends a message to Bob in order to compute the value of a function $f : \{0,1\}^n \times \{0,1\}^m \to \{0,1\}$. Essentially, Alice communicates a quantum state and Bob performs a measurement, both depending on their respective inputs. Though deceptively simple, this scenario is not at all fully understood. Let us just mention the following open problem: what is the largest complexity gap between quantum and classical protocols of this kind for computing a total Boolean function? The largest gap known is a factor of 2, as shown by Winter [Win04], but for all we know there could be examples where the gap is exponential, as it indeed is for certain partial functions [GKK$^+$08].

An interesting bound on such speedups can be found by investigating the effect of replacing quantum by classical messages. Let us sketch the proof of such a result. Suppose a total Boolean function f has a quantum one-way protocol with communication c, namely Alice sends c qubits to Bob, who can decide f with error $1/3$ by measuring Alice's message. We allow Alice and Bob to share an arbitrary input-independent entangled state. Extending Nayak's random access code bound [Nay99] Klauck [Kla00] showed that $Q^{A\to B,*}(f) \geq \Omega(VC(f))$, where $Q^{A\to B,*}(f)$ denotes the entanglement-assisted quantum one-way complexity of f, and $VC(f)$ the Vapnik-Chervonenkis dimension of the communication matrix of f. Together with Sauer's Lemma [Sau72] this implies that $D^{A\to B}(f) \leq O(Q^{A\to B,*}(f) \cdot m)$, where m is the length of Bob's input. See also [JZ09] for a related result.

A result such as the above is much more interesting in the case of partial functions. The reason is that for total functions a slightly weaker statement follows from a weak version of the random access code bound, which can be (and indeed has been [ANTV99]) established by the following argument: boost the quantum protocol for f until the error is below 2^{-2m}, where m is Bob's input length. Measure the message sent by Alice with *all* the measurements corresponding to Bob's inputs (this can be done with small total error) in order to determine Alice's row of the communication matrix and hence her input. This is a hard task by standard information theory facts (Holevo's bound). When considering partial functions a disaster happens: Bob does not know for which of his inputs y the value $f(x,y)$ is defined. If Bob measures the message for x with the observable for y and $f(x,y)$ is undefined any acceptance probability is possible and the message state can be destroyed.

Aaronson [Aar05] circumvented this problem in the following way: Bob now tries to learn Alice's message. He starts with a guess (the totally mixed state) and keeps a classical description of his guess. Alice also always knows what Bob's guess is. Bob can simulate quantum measurements by brute-force calculation: for any measurement operator Bob can simply calculate the result from his classical description. Alice can do the same. Since Bob has

some 2^m measurements he is possibly interested in, Alice can just tell him on which of these he will be wrong. Bob can then adjust his quantum state accordingly, and Aaronson's main argument is that he does not have to do this too often before he reaches an approximation of the message state. Note that Bob might never learn the message state if it so happens that all measurements are approximately correct on his guess. But if he makes a certain number of adjustments he will learn the message state and no further adjustments are needed.

Let us state Aaronson's result from [Aar05].

Fact 1.1. $D^{A\to B}(f) \leq O(Q^{A\to B}(f) \cdot \log(Q^{A\to B}(f)) \cdot m)$ for all partial Boolean $f : \{0,1\}^n \times \{0,1\}^m \to \{0,1\}$.

Aaronson later proved the following result, that removes the log-factor at the expense of having randomized complexity on the left hand side.

Fact 1.2. $R^{A\to B}(f) \leq O(Q^{A\to B}(f) \cdot m)$ for all partial Boolean $f : \{0,1\}^n \times \{0,1\}^m \to \{0,1\}$.

Our first result is the following improvement.

Result 1. $D^{A\to B}(f) \leq O(Q^{A\to B,*}(f) \cdot m)$ for all partial Boolean $f : \{0,1\}^n \times \{0,1\}^m \to \{0,1\}$.

Hence we remove the log-factor, and we allow the quantum communication complexity on the right hand side to feature prior entanglement between Alice and Bob. Arguably, looking into the entanglement assisted case (which is interesting for our second main result) led us to consider a more systematic progress measure than in Aaronson's proof, which in turn allowed us to analyze a different update rule for Bob that also works for protocols with error 1/3, instead of extremely small error as used in [Aar05], which is the cause of the lost log-factor.

We note that this result can be used to slightly improve on the "quantum-classical" simultaneous message passing lower bound for the Equality function by Gavinsky et al. [GRdW08], establishing a tight $\Omega(\sqrt{n})$ lower bound on the complexity of Equality in a model where quantum Alice and classical Bob (who do not share a public coin or entanglement) each send messages to the referee. The tight lower bound has also recently been established via a completely different and simpler method [GK14] (as well as generalized to a nondeterministic setting). Our result (as well as the one in [GK14]) allows a generalization to a slightly stronger model: Alice and the referee may share entanglement.

1.2 The Power of Quantum Proofs

We now turn to the second result of our paper, which is philosophically the more interesting. Interactive proof systems are a fundamental concept in computer science. Quantum proofs have a number of disadvantages: reading them may destroy them, errors may occur during verification, verification needs some sort of quantum machine, and it may be much harder to provide them than classical proofs. The main hope is that quantum proofs can in some situations be verified using fewer resources than classical proofs. Until now such a hope has not been verified formally. In the fully interactive setting Jain et al. have shown that the set of languages recognizable in polynomial time with the help of a quantum prover is equal to the set where the prover and verifier are classical (i.e., IP=QIP [JJUW11]).

The question remains open in the noninteractive setting. A question first asked by Aharonov and Naveh [AN02] and meriting much attention, is whether proofs that are quantum states can ever be easier to verify than classical proofs (by quantum machines) in the absence of interaction, i.e., whether the class QMA is larger than its analogue with classical proofs but quantum verifiers, known as QCMA. An indication that quantum proofs may be

powerful was given by Watrous [Wat00], who described an efficient QMA black box algorithm for deciding nonmembership in a subgroup. However, Aaronson and Kuperberg [AK07] later showed how to solve the same problem efficiently using a classical witness, giving a QCMA black box algorithm for the problem. They also introduced a quantum problem, for which they show that QMA black box algorithms are more efficient than QCMA black box algorithms. Using a quantum problem to show hardness for algorithms using classical proofs seems unfair though and a similar separation has remained open for Boolean problems.

In our second main result we compare the two modes of noninteractive proofs and quantum verification for a Boolean function in the setting of one-way communication complexity. More precisely we exhibit a partial Boolean function f, such that the following holds. f can be computed in a protocol where a prover who knows x, y can provide a quantum proof to Alice, and Alice sends quantum message to Bob, such that the total message length (proof plus message Alice to Bob) is $O(\log n)$. In the setting where a prover Merlin (still knowing all inputs) sends a classical proof to Alice, who sends a quantum message to Bob, the total communication is $\Omega(\sqrt{n}/\log n)$.

Result 2. There is a partial Boolean function f such that $QMA^{A\to B}(f) = O(\log n)$, while $QCMA^{A\to B,*}(f) = \Omega(\sqrt{n}/\log n)$.

We note that this is the first known exponential gap between computing Boolean functions in a QCMA and a QMA mode in any model of computation. Also, the lower bound is not too far from being tight, since there is an obvious upper bound of $O(\sqrt{n}\log n)$ for the problem.

So where does the power of quantum proofs come from in our result? Raz and Shpilka [RS04] show that QMA one-way protocols are as powerful as QMA two-way protocols. Their proof uses a quantum witness that is a superposition over the messages of different rounds. We show that for a certain problem with an efficient QMA protocol there is no efficient one-way QCMA protocol. Hence the weakness of classical proofs here is the impossibility of compressing interaction as in the QMA case.

2 Preliminaries

2.1 Quantum

For basic quantum background we refer to [NC00].

2.2 Communication Complexity Models

We assume familiarity with communication complexity, referring to [KN97] for more details about classical communication complexity and [Wol02] for quantum communication complexity.

For a partial Boolean function $f : \{0,1\}^n \times \{0,1\}^m \times \{0,1,\perp\}$, where $\perp$ stands for "undefined" the communication matrix A_f has rows labeled by $x \in \{0,1\}^n$ and columns labeled by $y \in \{0,1\}^m$, and entries $f(x,y)$. A protocol for f is correct, if it gives the correct output for all x, y with $f(x,y) \neq \perp$ (with certainty for deterministic protocols, and with probability $2/3$ for quantum protocols). A protocol is one-way, if Alice sends a message to Bob, who computes the function value, or vice versa. We denote by $D^{A\to B}(f)$ the deterministic one-way communication complexity of a function f, when Alice sends the message to Bob.

Two rows x, x' of A_f are *distinct*, if there is a column y, such that $f(x,y) = 1$ and $f(x',y) = 0$ or vice versa, i.e., the function values differ on some defined input. Note that being not distinct is not an equivalence relation: x, x' can be not distinct, as well as x, x'', while x, x' are distinct. Nevertheless a one-way protocol for f needs to group inputs x into messages such that no two distinct x, x' share the same message.[1]

Similar to the above, we denote by $Q^{A\to B}(f)$ the quantum one-way communication complexity of f with error 1/3. This notion is of course asymptotically robust when it comes to changing the error to any other constant. $Q^{A\to B,*}(f)$ denotes the complexity if Alice and Bob share entanglement.

We now define some more esoteric modes of communication that extend the standard nondeterministic mode to the quantum case. We restrict our attention to one-way protocols.

Definition 1. *In a one-way MA-protocol there are 3 players Merlin, Alice, Bob. Merlin sends a classical message to Alice, who sends a classical message to Bob, who gives the output. Alice and Bob share a public coin, which is not seen by Merlin. For a Boolean function $f : \{0,1\}^n \times \{0,1\}^m \to \{0,1\}$ the protocol is correct, if for all 1-inputs there is a message from Merlin, such that with probability 2/3 Bob will accept, whereas for all 0-inputs, and all messages from Merlin, Bob will reject with probability 2/3. The communication complexity is defined as usual and denoted by $MA^{A\to B}(f)$.*

A one-way QCMA-protocol is defined similarly, but whereas Merlin's message is still classical, Alice can send a quantum message to Bob, and Alice and Bob may share entanglement. The complexity with shared entanglement is denoted $QCMA^{A\to B,*}(f)$.

In a one-way QMA-protocol also Merlin's message may be quantum. The complexity is denoted by $QMA^{A\to B}(f)$ in the case where no entanglement is allowed.

2.3 Quantum Information Measures

In this paper we need only a few well established notions of information and distinguishability. A density matrix is a positive semidefinite matrix of trace 1. Density matrices will also be referred to as quantum states in this paper.

Definition 2. *The von Neumann entropy of a quantum state ρ is $S(\rho) = -\mathsf{Tr}\rho\log\rho$.*

The relative von Neumann entropy of quantum states ρ, σ is $S(\rho||\sigma) = \mathsf{Tr}\rho\log\rho - \mathsf{Tr}\rho\log\sigma$ if supp $\rho \subseteq$ supp σ, otherwise $S(\rho||\sigma) = \infty$.

The relative min-entropy of ρ, σ is $S_\infty(\rho||\sigma) = \inf\{c : \sigma - \rho/2^c$ is positive semidefinite$\}$.

It is easy to see that $S(\rho||\sigma) \le S_\infty(\rho||\sigma)$, see [Dat09] for a proof. An important measure of how far apart quantum states are is the trace distance.

Definition 3. *The trace norm of a Hermitian operator ρ is defined as $||\rho||_t = \mathsf{Tr}\sqrt{\rho\rho^\dagger}$.*

The trace distance between ρ and σ is $||\rho - \sigma||_t$.

We list two well known facts. First, Uhlmann monotonicity.

Fact 2.1. If $\tilde{\rho}, \tilde{\sigma}$ result from measuring ρ, σ then $S(\tilde{\rho}||\tilde{\sigma}) \le S(\rho||\sigma)$.

[1] It is instructive to consider the function $f(x,i;y,j) = x_{i\oplus j}$ under the promise that $x = y$. This function has only n *distinct* rows and columns, and $D^{A\to B}(f) = D^{B\to A}(f) \le O(\log n)$. Nevertheless A_f has many more actual rows and columns. Trying to reduce the number of actual columns to a set of distinct columns increases the number of distinct rows. Hence one has to be careful when considering partial functions.

Secondly, the quantum Pinsker inequality [HOT81], see also [KNTZ01].

Fact 2.2. $||\rho - \sigma||_t \leq \sqrt{2 \ln 2 \; S(\rho||\sigma)}$.

We note that any two states that are close in trace distance are hard to distinguish by any measurement, namely the classical distance between the measurement results is at most the trace distance of the measured states.

3 Making Quantum Messages Deterministic

Theorem 3.1. *For every partial Boolean function* $f : \{0,1\}^n \times \{0,1\}^m \to \{0,1\}$ *we have* $D^{A \to B}(f) \leq O(Q^{A \to B,*}(f) \cdot m)$.

We note that in the case of total functions the theorem follows from a result in [Kla00] combined with Sauer's lemma [Sau72], and that two weaker versions of the theorem have been proved by Aaronson: in [Aar05] he shows the result with an additional log-factor on the right hand side, and without allowing entanglement, in [Aar07] with $R^{A \to B}(f)$ on the left hand side (and no additional log-factor), but again without entanglement.

Our proof follows Aaronson's main approach in [Aar05], in which Bob maintains a classical description of a quantum state as his guess for the message he should have received, and Alice informs him about inputs on which this state will perform badly, so that he can adjust his guess. His goal is to either get all measurement results approximately right, or to learn the message state. We will refer to these states as the current guess state, and the target state.

We deviate from Aaronson's proof in two ways. First, we work with a different progress measure that is more transparent than Aaronson's, namely the relative entropy between the target state and the current guess. This already allows us to work in the entanglement-assisted case.

Secondly, we modify the rule by which Bob updates his guess. In Aaronson's proof Bob projects his guess state onto the subspace on which the target state has a large projection (because the message is accepted by the corresponding measurement with high probability). This has the drawback that one cannot use the actual message state of the protocol as the target state, because that state usually has considerable projection onto the orthogonal complement of the subspace, making the relative entropy infinitely large! Hence Aaronson uses a boosted and projected message state as the target state. This state is close to the actual message state thanks to the boosting, and projection of the guess state now properly decreases the relative entropy, since the target state is fully inside the subspaces. The boosting step costs exactly the log-factor we aim to remove.

So in the situation where Bob wants to update his guess state σ, knowing that the target state ρ will be accepted with probability $1 - \epsilon$ when measuring the observable consisting of subspace V_y and its complement, we let Bob replace σ with the mixture of $1 - \epsilon$ times the projection onto V_y and ϵ times the projection onto $V_y^\perp$. The main part of the proof is then to show that this decreases the relative entropy $S(\rho||\sigma)$ given that $\mathsf{Tr}(V_y \sigma) < 1 - 10\sqrt{\epsilon}$, i.e., in case σ was not good enough already. Eventually either all measurements can be done by Bob giving the correct result, or the current guess state σ satisfies $S(\rho||\sigma) \leq 5\sqrt{\epsilon}$, in which case ρ and σ are also close in the trace distance meaning that any future measurement will give almost the same results on both states.

Proof. Fix any entanglement assisted one-way protocol with quantum communication $q = Q^{A\to B,*}(f)$. Using standard boosting we may assume that the error of the protocol is at most $\epsilon = 10^{-6}$ for any input x, y. This increases the communication by a small constant factor at most.

Using teleportation we can replace the quantum communication by $2q$ classical bits of communication at the expense of adding q EPR-pairs to the shared entangled state. Let $|\phi\rangle$ denote the entangled state shared by the new protocol. We can assume this is a pure state, because if this is not the case we may consider any purification, and Alice and Bob can ignore the purification part. Note that we do not restrict the number of qubits used in $|\phi\rangle$.

In the protocol, for a given input x Alice has to perform a unitary transformation on her part of $|\phi\rangle$ (we assume that any extra space used is also included in $|\phi\rangle$ and that measurements are replaced by unitaries) and then sends a classical message. Bob first applies the unitary from the teleportation protocol (which only depends on the classical message). Let's denote the state shared by Alice and Bob at this point by $|\phi_x\rangle$. Following this Bob performs a measurement (depending on his input y) on his part of $|\phi_x\rangle$. This measurement determines the output of the protocol on x, y. We may assume by standard techniques that Bob's measurements are projection measurements, and that the subspaces used in the projection measurements have dimension $d/2$, where d is the dimension of the underlying Hilbert space.

Recall that $|\phi\rangle$ and $|\phi_x\rangle$ are bipartite states shared by Alice and Bob. Let $\rho = \mathsf{Tr}_A|\phi_x\rangle\langle\phi_x|$ and $\sigma_1 = \mathsf{Tr}_A|\phi\rangle\langle\phi|$, i.e., the states when Alice's part is traced out. Bob wants to learn ρ in order to be able to determine all measurement results on ρ. We show how to do this with $O(m \cdot q)$ bits of deterministic communication from Alice. Note that the state σ_1 is known to Bob in the sense that he knows its classical description.

Since Alice's local operations do not change Bob's part of $|\phi\rangle\langle\phi|$, the difference between ρ and σ_1 is introduced via the correction operations in the teleportation protocol that Bob applies after he receives Alice's message. But with probability 2^{-2q} Bob does not have to do anything, i.e., when Alice's message is the all 0-s string. This implies that

$$\sigma_1 = \frac{1}{2^{2q}}\rho + \theta,$$

for some positive semidefinite θ with trace $1 - 1/2^{2q}$. Hence we get that

$$S(\rho||\sigma_1) \leq S_\infty(\rho||\sigma_1) \leq 2q.$$

In other words, Bobs target ρ and initial guess σ_1 have small relative entropy.

We can now describe the protocol. Bob starts with the classical description of σ_1. This state is also known by Alice, since it does not depend on the input. Throughout the protocol Bob will hold states σ_i, which will be updated when needed, using information provided by Alice. Bob also has a set of measurement operators $P_y, I - P_y$ for all his inputs y. Bob and Alice each loop over his inputs y, and compute $p_y = \mathsf{Tr}(P_y\sigma_i)$. This is the acceptance probability, if σ_i is measured with the measurement for his input y. Alice also computes $p'_y = \mathsf{Tr}(P_y\rho)$ which is the acceptance probability of the quantum protocol. If p_y and p'_y are too far apart, Alice will notify Bob of the correct acceptance probability on y (with precision ϵ^2), which takes $m + O(1)$ bits of communication.

Alice does not send a message if $f(x, y)$ is undefined, because the acceptance probability on such inputs is irrelevant. Suppose $p'_y = 1 - \epsilon_y$, where $\epsilon_y \leq \epsilon$ is the error on x, y (and $f(x, y) = 1$), but $p_y = 1 - a$ for some $1 > a \geq 10\sqrt{\epsilon}$. If this is not the case the measurement

for y applied to σ_i already yields the correct result and no information from Alice is needed. So if p_y, p'_y are far apart Alice will send y (using m bits) and ϵ_y as a floating point number with precision ϵ^2 (using $O(1)$ bits).

Bob then adjusts σ_i to obtain a state σ_{i+1}. Suppose he knows the correct ϵ_y (the difference between ϵ_y and its approximation sent by Alice will be irrelevant). This means that $\mathsf{Tr}(P_y\sigma_i) = 1-a$ but $\mathsf{Tr}(P_y\rho) = 1-\epsilon_y$. P_y is the projector onto a subspace V_y. We have assumed that each V_y has dimension $d/2$ if d is the dimension of the underlying Hilbert space. Let B_i denote an orthonormal basis, in which the first $d/2$ elements span V_y, and the remaining $d/2$ span $V_y^\perp$. Furthermore in this basis the upper left and lower right quadrants of σ_i are diagonal. Hence σ_i will look like this:

$$\sigma_i = \begin{pmatrix} A & B \\ B^* & D \end{pmatrix} = \begin{pmatrix} * & 0 & 0 & * & \cdots & * \\ 0 & \ddots & 0 & * & \vdots & * \\ 0 & 0 & * & * & \cdots & * \\ * & \cdots & * & * & 0 & 0 \\ * & \vdots & * & 0 & \ddots & 0 \\ * & \cdots & * & 0 & 0 & * \end{pmatrix},$$

where $\mathsf{Tr}(A) = 1-a$ and $\mathsf{Tr}(D) = a$. We can now define

$$\sigma_{i+1} = \begin{pmatrix} \frac{1-\epsilon_y}{1-a}A & 0 \\ 0 & \frac{\epsilon_y}{a}D \end{pmatrix},$$

σ_{i+1} is diagonal in our basis B_i. Clearly σ_{i+1} would perform exactly as desired on measurement $P_y, I-P_y$. But Bob already knows the function value on y and can carry on with the next y.

Before we continue we have to argue that the case $a = 1$ can never happen. Since $\mathsf{Tr}(P_y\rho) = 1-\epsilon_y > 0$ the state ρ has a nonzero projection onto V_y. If $a = 1$ then σ_i sits entirely in $V_y^\perp$, and hence $S(\rho||\sigma_i) = \infty$. But since we start with a finite $S(\rho||\sigma_1)$ and only decrease that value the situation $a = 1$ is impossible.

Coming back to the protocol, it is obvious that Bob will learn the correct value of $f(x,y)$ for all y such that $f(x,y)$ is defined. Hence the protocol is deterministic and correct. The remaining question is how many times Alice has to send a message to Bob. We will show that this happens at most $O(Q^{A\to B,*}(f)/\sqrt{\epsilon})$ times, which establishes our theorem.

The main claim that remains to be shown is the following.

Claim 3.2. $S(\rho||\sigma_i) \geq S(\rho||\sigma_{i+1}) + a/2$ if $a \geq 10\sqrt{\epsilon}$.

This establishes the upper bound on the number of messages, because the relative entropy, which starts at $2q$ is decreased by $a/2 \geq 5\sqrt{\epsilon}$ for each message. After at most $2q/(5\sqrt{\epsilon})$ iterations the protocol has either ended (in which case Bob might never learn ρ, but will still know all measurement results), or we have

$$S(\rho||\sigma_T) \leq 5\sqrt{\epsilon}.$$

To see this assume we are still in the situation of the claim. The claim states that the relative entropy can be reduced by $a/2$ as long as $a \geq 10\sqrt{\epsilon}$. So the process stops (assuming we don't run out of suitable $y's$) no earlier than when $S(\rho||\sigma_i) < a/2 \leq 5\sqrt{\epsilon}$.

But then by the quantum Pinsker inequality we have that at the final time T: $||\rho - \sigma_T||_t \leq \sqrt{10 \ln 2\sqrt{\epsilon}} < 0.1$ in the end, and hence for *all* measurements their results are close. Hence no more than $O(q/\sqrt{\epsilon}) = O(q)$ messages have to be sent.

We now finish the proof by showing the claim.

Define another state

$$\tilde{\sigma_i} = \begin{pmatrix} A & 0 \\ 0 & D \end{pmatrix},$$

i.e., the state σ_i with its upper right and lower left quadrants deleted. It is easy to see that this is still a density matrix. Indeed $\tilde{\sigma_i}$ is the state σ_i after measuring $P_y, I - P_y$. Also define $\tilde{\rho}$ to be the matrix ρ with its upper left and lower right quadrants replaced by 0, which is again a density matrix resulting from measuring ρ. If

$$\rho = \begin{pmatrix} E & F \\ G & H \end{pmatrix} \text{ then } \tilde{\rho} = \begin{pmatrix} E & 0 \\ 0 & H \end{pmatrix}.$$

By Uhlmann monotonicity we have $S(\rho||\sigma_i) \geq S(\tilde{\rho}||\tilde{\sigma_i})$, and it suffices to show that $S(\tilde{\rho}||\tilde{\sigma_i}) \geq S(\rho||\sigma_{i+1}) + a/2$. Note that both $\tilde{\sigma_i}$ and σ_{i+1} are diagonal in the basis B_i. Furthermore $S(\rho||\sigma_{i+1}) = S(\tilde{\rho}||\sigma_{i+1}) + S(\tilde{\rho}) - S(\rho)$.

We first bound the term $S(\tilde{\rho}) - S(\rho)$. In the basis B_i we may view ρ as a bipartite state ρ_{RQ} consisting of a qubit R (corresponding to membership in V_y) and the remaining qubits Q. In $\tilde{\rho}$ the qubit R has been measured. Consider attaching another qubit T, and instead of measuring R applying the unitary that "copies" R to T. After the unitary $S(\rho) = S(\rho_{QRT}) \geq S(\rho_{QR}) - S(\rho_T) = S(\tilde{\rho}) - S(\rho_T)$, due to the Araki-Lieb inequality and because $\rho_{QR} = \tilde{\rho}$. But $S(\rho_T) = H(\epsilon_y)$ and hence $S(\tilde{\rho}) - S(\rho) \leq H(\epsilon_y)$.

We need to compare $\mathsf{Tr}(\tilde{\rho} \log \tilde{\sigma_i})$ and $\mathsf{Tr}(\tilde{\rho} \log \sigma_{i+1})$. Note that $\tilde{\sigma_i}$ and σ_{i+1} are both diagonal.

$$\begin{aligned}
& \mathsf{Tr}(\tilde{\rho} \log \tilde{\sigma_i}) \\
= & \sum_j \tilde{\rho}(j,j) \log(\tilde{\sigma_i}(j,j)) \\
= & \sum_{j \leq d/2} \tilde{\rho}(j,j) \log(\sigma_{i+1}(j,j) \cdot \frac{1-a}{1-\epsilon_y}) \\
+ & \sum_{j > d/2} \tilde{\rho}(j,j) \log(\sigma_{i+1}(j,j) \cdot \frac{a}{\epsilon_y}) \\
= & \sum_j \tilde{\rho}(j,j) \log(\sigma_{i+1}(j,j)) + (1-\epsilon_y) \cdot \log \frac{1-a}{1-\epsilon_y} + \epsilon_y \cdot \log \frac{a}{\epsilon_y} \\
= & \mathsf{Tr}(\tilde{\rho} \log \sigma_{i+1}) + H(\epsilon_y) - H(\epsilon_y, a),
\end{aligned}$$

where $H(u,v) = -u \log v - (1-u) \log(1-v)$.

Assuming that $a \geq 10\sqrt{\epsilon}$ and $\epsilon_y \leq \epsilon = 10^{-6}$ we can estimate

$$H(\epsilon_y, a) \geq a. \tag{1}$$

If $a > 1/2$ then (1) is true from the first term $(1-\epsilon_y)\log(1/(1-a)) \geq a$. Otherwise for all $a \in [0,1]$ we have $-\log(1-a) \leq a$, hence the first term in $H(\epsilon_y, a)$ is at least $(1-\epsilon_y)a$. And for the second term $\epsilon_y \log(1/a) \geq \epsilon_y \log(2) \geq \epsilon_y$, so (1) is always true.

Also note that $H(\epsilon_y) \leq H(\epsilon) \leq \epsilon \log(\epsilon^{-1}) \leq \sqrt{\epsilon} \leq a/10$.
But then

$$\begin{aligned}
& S(\rho||\sigma_i) \\
\geq\ & S(\tilde{\rho}||\tilde{\sigma_i}) \\
=\ & -S(\tilde{\rho}) - \mathsf{Tr}(\tilde{\rho} \log \tilde{\sigma_i}) \\
\geq\ & -S(\rho) - H(\epsilon) - \mathsf{Tr}(\tilde{\rho} \log \tilde{\sigma_i}) \\
\geq\ & -S(\rho) - H(\epsilon) - \mathsf{Tr}(\tilde{\rho} \log \sigma_{i+1}) + a - H(\epsilon) \\
=\ & -S(\rho) - \mathsf{Tr}(\rho \log \sigma_{i+1}) + a - 2H(\epsilon) \\
\geq\ & S(\rho||\sigma_{i+1}) + a/2.
\end{aligned}$$

□

4 Quantum versus Classical Proofs

Let us first define the problem for which we prove our separation result.

Definition 4. *The function* $\mathsf{MajIx}(x, I)$*, where* $I = \{i_1, \ldots, i_{\sqrt{n}}\}$*, each* $i_j \in \{1, \ldots, n\}$*, and* $x \in \{0,1\}^n$ *is defined as follows:*

1. *if* $|\{j : x_{i_j} = 1\}| = \sqrt{n}$ *then* $\mathsf{MajIx}(x, I) = 1$,
2. *if* $|\{j : x_{i_j} = 1\}| \leq 0.9\sqrt{n}$ *then* $\mathsf{MajIx}(x, I) = 0$,
3. *otherwise* $\mathsf{MajIx}(x, I)$ *is undefined.*

The function has been studied in [Kla11], where it is shown that one-way MA protocols for the problem need communication $\Omega(\sqrt{n})$. Our main technical result here is to extend this to one-way QCMA protocols.

It is obvious on the other hand, that there is a cheap protocol when Bob can send a message to Alice.

Lemma 4.1. $R^{B\to A}(\mathsf{MajIx}) = O(\log n)$.

Raz and Shpilka [RS04] show that any problem with $QMA(f) = c$ (i.e., QMA protocol where Alice and Bob can interact over many rounds) has a QMA protocol of cost $poly(c)$ in which Merlin sends a message to Alice, who sends a message to Bob. By inspection of their proof the polynomial overhead can be removed in the case of constant rounds of interaction between Alice and Bob.

Lemma 4.2. *If* $QMA(f) = c$ *and this cost can be achieved by a protocol with* $O(1)$ *rounds, then* $QMA^{A\to B}(f) = O(c)$.

We give more details in Appendix A. The lemma immediately implies the following.

Theorem 4.3. $QMA^{A\to B}(\mathsf{MajIx}) = O(\log n)$.

We give a self-contained proof of this fact in Appendix B. Our protocol has completeness 1, hence even the one-sided error version of $QMA^{A\to B}$ is separated from $QCMA^{A\to B}$ by the following lower bound.

Theorem 4.4. $QCMA^{A\to B,*}(\mathsf{MajIx}) \geq \Omega(\sqrt{n}/\log n)$.

Hence we can conclude the following.

Corollary 4.5. *There is a partial Boolean function f such that $QMA^{A\to B}(f) = O(\log n)$, while $QCMA^{A\to B,*}(f) = \Omega(\sqrt{n}/\log n)$.*

Proof. Fix any QCMA protocol $\mathcal{P}$ for MajIx. Furthermore define a distribution on inputs as follows. Fix any error correcting code $C \subseteq \{0,1\}^n$ with distance $n/4$ (i.e., every two codewords have Hamming distance at least $n/4$). Such codes of size $2^{\Omega(n)}$ exist by the Gilbert-Varshamov bound. We do not care about the complexity of decoding and encoding for our code. Furthermore we require the code to be balanced, i.e., that any codeword has exactly $n/2$ ones. This can also be achieved within the stated size bound. For our distribution on inputs first choose $x \in C$ uniformly, and then uniformly choose I among all subsets of $\{1,\ldots,n\}$ of size $\sqrt{n}$. Note that the probability of 1-inputs under the distribution μ just defined is between $2^{-\sqrt{n}}$ and $2^{-\sqrt{n}-1}$ due to the balance condition on the code.

If the cost (i.e., communication from Merlin plus communication from Alice) of $\mathcal{P}$ is c, then there are at most 2^c different classical proofs sent by Merlin. We identify such proofs p with the set of 1-inputs that are accepted by the protocol with probability at least $2/3$ when using the proof p. Hence there must be a proof P containing 1-inputs of measure at least $2^{-\sqrt{n}-c-1}$, because for every 1-input there is a proof with which it is accepted with probability $2/3$ or more. Furthermore, given P, no 0-input is accepted with probability larger than $1/3$. Note that inputs outside of the promise, or 1-inputs outside of P can be accepted with any probability between 0 and 1. Denote by f_P the partial function $\{0,1\}^n \times \{0,1\}^m \to \{0,1,\perp\}$, in which all inputs in P are accepted, and all 0-inputs of f are rejected, and the remaining inputs have undefined function value ($\perp$). $m = \Theta(\sqrt{n}\log n)$ is the length of Bob's input. Denote by $M = \binom{n}{\sqrt{n}}$ the number of Bob's inputs.

Obviously f_P can be computed by a one-way quantum protocol without prover using communication c (and possibly using shared entanglement between Alice and Bob). Now due to Theorem 3.1 this implies that $D^{A\to B}(f_P) \leq O(c\cdot m)$. We will argue that on the other hand $D^{A\to B}(f_P) \geq \Omega(n)$, and hence $c \geq \Omega(n/m) = \Omega(\sqrt{n}/\log n)$, which is our theorem.

Denote by A the communication matrix of f_P. A row of A is *fat*, if it contains more than $M2^{-\sqrt{n}-2c}$ 1-inputs (to f_P). Note that there can be no fewer than $|C|2^{-c-2}$ fat rows, because there are at least $|C|M2^{-\sqrt{n}-c-1}$ 1-inputs in P, the non-fat rows contain together at most $|C|M2^{-\sqrt{n}-2c}$ 1-inputs and each fat row at most $M2^{-\sqrt{n}}$ 1-inputs. Let C' denote the row set consisting of the fat rows only, and A' the matrix A restricted to those rows. We claim that A' has C' distinct rows. Recall that for distinct rows there is a column, where one row has a 1 entry and the other a 0 entry. This means that $D^{A\to B}(f_P) \geq \log C' \geq \log|C| - c - 2 \geq \Omega(n)$ (unless $c = \Omega(n)$ already).

To show that all pairs of rows in A' are distinct consider two of them, named x, y. We identify the row labeled by x with the set for which x is the characteristic vector. Recall x, y are both codewords and are both fat. $x \cap y \leq n(1/2 - 1/8)$ because x and y have Hamming distance at least $n/4$. Let $S \subseteq \mathcal{I}$ be the set of I such that $(x, I) \in P$. Then for all $I \in S$ we have that all of the $i \in I$ must satisfy $x_i = 1$.

Furthermore let $T \subseteq S$ be the set $I \in S$, such that $|I \cap x \cap y| \geq 0.9\sqrt{n}$. Then

$$|T| \leq \sqrt{n} \cdot \binom{\frac{3n}{8}}{0.9\sqrt{n}} \cdot \binom{\frac{n}{8}}{0.1\sqrt{n}} \leq 2^{-\alpha\sqrt{n}} \cdot \binom{\frac{n}{2}}{\sqrt{n}},$$

for some constant $\alpha > 0$. Note that the binomial coefficient on the right hand side is the number of $I \in \mathcal{I}$ such that (x, I) is a 1-input. Hence

$$\frac{\mu((\{x\} \times T) \cap P)}{\sum_{I \in \mathcal{I}: \mathsf{MajIx}(x,I)=1} \mu(x, I)} \leq 2^{-\alpha\sqrt{n}},$$

i.e., a small fraction of 1-inputs (x, I) in P on row x have $I \in T$, but x is fat and has more 1-inputs. We can assume that $c < \alpha\sqrt{n}/10$ and hence $2^{-2c} \geq 2 \cdot 2^{-\alpha\sqrt{n}}$, so that the set T contributes little to the set of $I \in \mathcal{I}$ with $(x, I) \in P$. In particular there must be at least one $I \notin T$ such that (x, I) is in P, and hence $f_P(x, I) = 1$.

So let us examine the set of all $I \in S - T$ (so $|I \cap x \cap y| < 0.9\sqrt{n}$). $|I \cap x| = \sqrt{n}$ and hence $|y \cap I| \leq .9\sqrt{n}$, i.e., (y, I) is a 0-input. This gives us the desired column I, such that $f_P(x, I) = 1$ and $f_P(y, I) = 0$, i.e., x, y are two distinct rows in A'.

□

5 Open Problems

- Aaronson [Aar07] argues that the bound in Theorem 3.1 is tight to within polylogarithmic factors for partial functions. However, a longstanding conjecture is that for all total Boolean functions f we have $R^{A\to B}(f) = O(Q^{A\to B}(f))$, or something weaker, i.e., that for total Boolean functions we can remove the factor m completely (possibly at the expense of increasing the dependence on $Q^{A\to B}(f)$ polynomially). An easier problem might be to replace m by something smaller like $\sqrt{m}$ for total functions.

- For many "nondeterministic" modes of communication complexity one-way communication is as good as two-way communication, for instance for nondeterministic, QMA, AM-complexity. We have proved that this is not the case for QCMA protocols (and this was known previously for MA-protocols [Kla11]). Is there a proper round-hierarchy for QCMA or MA protocols, i.e, is it true that there is a function that can be computed efficiently in k rounds but not in $k-1$ rounds?

- MA-communication complexity has recently been applied to the analysis of cloud-computing on data-streams and related topics [CCM+13, CCM09]. Currently we don't have lower bounds larger than $\sqrt{n}$ for MA-communication complexity of explicit functions, while counting arguments show that most functions have complexity $\Omega(n)$. This gap is quite significant in practice. Can larger bounds be shown for an explicit function, at least in the one-way model, or the even more restricted *online* one-way model?

- Lower bounds for the AM-communication complexity of any explicit function remain elusive.

- It would be interesting to separate QCMA- and QMA- communication complexity in the general two-way communication model. Such a separation could be used in the algebrization framework [AW09] to argue that showing QCMA=QMA (in the Turing machine world) would require nonalgebrizing techniques, and might also be thought of as evidence that these classes are not equal after all.

- A similar problem, and maybe less ambitious, is to show that the QCMA- and QMA-query complexities of a Boolean problem are very different. However, while for communication complexity we have a good candidate for such a separation (the QMA-complete problem), we are not aware of a good candidate for the query complexity setting. Is there a complete (promise) problem for QMA-query complexity?

Acknowledgement

The authors thank Dmitry Gavinsky for helpful discussions.

References

[Aar05] S. Aaronson. Limitations of quantum advice and one-way communication. *Theory of Computing*, 1:1–28, 2005. Earlier version in Complexity'04. quant-ph/0402095.

[Aar07] S. Aaronson. The learnability of quantum states. *Proceedings of the Royal Society of London*, A463(2088), 2007. quant-ph/0608142.

[AK07] Scott Aaronson and Greg Kuperberg. Quantum versus classical proofs and advice. *Theory of Computing*, 3(1):129–157, 2007.

[AN02] D. Aharonov and T. Naveh. Quantum NP - a survey. quant-ph/0210077, 2002.

[ANTV99] A. Ambainis, A. Nayak, A. Ta-Shma, and U. Vazirani. Dense quantum coding and a lower bound for 1-way quantum automata. In *Proceedings of 31st ACM STOC*, pages 697–704, 1999.

[AW09] Scott Aaronson and Avi Wigderson. Algebrization: A new barrier in complexity theory. *TOCT*, 1(1), 2009.

[CCM09] Amit Chakrabarti, Graham Cormode, and Andrew Mcgregor. Annotations in data streams. In *ICALP*, 2009.

[CCM+13] A. Chakrabarti, G. Cormode, A. McGregor, J. Thaler, and S. Venkatasubramanian. On interactivity in arthur-merlin communication and stream computation. ECCC TR13-180, 2013.

[Dat09] Nilanjana Datta. Min- and max- relative entropies and a new entanglement monotone. *IEEE Transactions on Information Theory*, 55:2816–2826, 2009.

[GK14] D. Gavinsky and H. Klauck. Equality, Revisited. Manuscript, 2014.

[GKK+08] Dmitry Gavinsky, Julia Kempe, Iordanis Kerenidis, Ran Raz, and Ronald de Wolf. Exponential separation for one-way quantum communication complexity, with applications to cryptography. *SIAM J. Comput.*, 38(5):1695–1708, 2008.

[GRdW08] D. Gavinsky, O. Regev, and R. de Wolf. Simultaneous Communication Protocols with Quantum and Classical Messages. *Chicago Journal of Theoretical Computer Science*, 7, 2008.

[HOT81] F. Hiai, M. Ohya, and M. Tsukada. Sufficiency, KMS condition and relative entropy in von Neumann algebras. *Pacific J. Math.*, 96:99–109, 1981.

[JJUW11] Rahul Jain, Zhengfeng Ji, Sarvagya Upadhyay, and John Watrous. QIP = PSPACE. *J. ACM*, 58(6), 2011.

[JZ09] R. Jain and S. Zhang. New bounds on classical and quantum one-way communication complexity. *Theoretical Computer Science*, 410(26):2463–2477, 2009.

[Kla00] H. Klauck. On quantum and probabilistic communication: Las Vegas and one-way protocols. In *Proceedings of 32nd ACM STOC*, pages 644–651, 2000.

[Kla11] H. Klauck. On Arthur Merlin Games in Communication Complexity. In *IEEE Conference on Computational Complexity*, pages 189–199, 2011.

[KN97] E. Kushilevitz and N. Nisan. *Communication Complexity*. Cambridge University Press, 1997.

[KNTZ01] H. Klauck, A. Nayak, A. Ta-Shma, and D. Zuckerman. Interaction in quantum communication complexity. In *Proceedings of 33rd ACM STOC*, 2001.

[Nay99] A. Nayak. Optimal lower bounds for quantum automata and random access codes. In *Proceedings of 40th IEEE FOCS*, pages 369–376, 1999. quant-ph/9904093.

[NC00] M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000.

[RS04] R. Raz and A. Shpilka. On the power of quantum proofs. In *Proceedings of Computational Complexity*, pages 260–274, 2004.

[Sau72] N. Sauer. On the density of families of sets. *J. Combin. Theory Ser. A*, 13:145–147, 1972.

[Wat00] J. Watrous. Succinct quantum proofs for properties of finite groups. In *Proceedings of 41st IEEE FOCS*, pages 537–546, 2000. quant-ph/0011023.

[Win04] A. Winter. Quantum and classical message identification via quantum channels. In *Festschrift A.S. Holevo 60*, pages 171–188, 2004.

[Wol02] R. de Wolf. Quantum communication and complexity. *Theoretical Computer Science*, 287(1):337–353, 2002.

A Proof Idea of Lemma 4.2

The proof of Raz and Shpilka in [RS04] proceeds by showing that a problem LSD is complete for QMA communication complexity. It is easy to see that LSD can be computed by a QMA one-way protocol with logarithmic communication. The main difficulty is showing that any problem f with $QMA(f) = C$ can be reduced to an instance of LSD of size $2^{poly(C)}$. Here we will argue that if the QMA complexity of f is C for protocols that have only $O(1)$ rounds between Alice and Bob, then the constructed LSD instance has size $2^{O(C)}$ only, and hence we get that $QMA^{A\to B}(f) = O(C)$.

For the problem LSD (linear space distance), Alice and Bob each receive a subspace of $\mathbb{R}^n$ of dimension $n/4$. The distance between two subspaces V, W is the minimum over all pairs of unit vectors from V and W of the euclidean distance between the vectors. For LSD the promise is that the spaces given to Alice and Bob are either very close (distance at most $0.1\sqrt{2}$) or very far (distance at least $0.9\sqrt{2}$). The 1-inputs of LSD are the pairs of spaces that are close. Note that there is a QMA one-way protocol for this problem of logarithmic cost: Merlin sends Alice a vector that is close to both subspaces as a quantum state, Alice measures the state with the observable consisting of her subspace and its complement, and if the state projects into her subspace she sends the projected state to Bob (otherwise she rejects), who measures with his observable. If both measurements succeed, they accept.

The main lemma in the reduction of [RS04] is the following (their Lemma 19):

Lemma A.1. *If f has a QMA protocol with proof length p, communication c, and r rounds and error $1/r^4$, then f can be reduced (by local operation by Alice and Bob) to an instance of LSD where the underlying space is $\mathbb{R}^{(r+1)2^{2(c+p)}}$, and the distance between the two spaces is at least $1/r^{1.5}$ in the case of 0-inputs, and at most $\sqrt{2}/r^{2.5}$ in the case of 1-inputs.*

This lemma is then combined with standard boosting ideas to improve the gap in the distance, but for us the fact that amplifying the success probability by parallel repetition in a QMA one-way protocol is possible suffices. Furthermore, in the case of $r = O(1)$, the LSD instance provided by the above lemma already has a constant gap and the correct size of $2^{O(c+p)}$. Reducing a function f to the LSD instance, and then running $O(1)$ times the protocol for LSD in parallel suffices to get a one-way QMA-protocol for f.

B Proof of Theorem 4.3

We describe a protocol with constant gap, which can be amplified to have the usual soundness using standard techniques.

Alice holds a string $x \in \{0,1\}^n$, Bob a set I of $\sqrt{n}$ distinct indices. Alice receives a proof from Merlin, which is supposed to be the uniform superposition $\sum_{i \in I} |i\rangle$. Alice attaches one more qubit, applies the unitary that maps $|i\rangle|a\rangle$ to $|i\rangle|x_i \oplus a\rangle$, then measures that extra qubit. If the result is 0 she rejects. Otherwise she discards the extra qubit and sends the remaining state to Bob. Bob measures this state with the observable that consists of the subspace spanned by the vector with 1 in all positions $i \in I$ and 0 elsewhere, and its orthogonal complement. He accepts, if this measurement projects onto the first subspace. Note that the communication is $\log n$ qubits from both Merlin and Alice.

If Merlin is honest (and $x_i = 1$ for all $i \in I$), he sends the uniform superposition $\sum_{i \in I} |i\rangle$. Alice's measurement will not change this state (she effectively projects into the space spanned by the $|i\rangle$ with $x_i = 1$). Bob's measurement will accept with certainty. Hence the protocol has completeness 1.

Now assume that at most $0.9\sqrt{n}$ of the $i \in I$ satisfy that $x_i = 1$. We may assume that Merlin's message is a pure state of the correct dimension. For each mixed state there is a pure state that will perform at least as well, and states with the wrong dimension will be rejected immediately. Again, unless Alice rejects, her measurement projects into the space spanned by $|i\rangle$ for which $x_i = 1$. Denote this state by $|\psi\rangle = \sum_{i:x_i=1} \alpha_i |i\rangle$. Bob measures the observable consisting of $\mathrm{span}(|\phi_I\rangle)$, where $|\phi_I\rangle = \sum_{i \in I} |i\rangle / n^{1/4}$, and its orthogonal complement. The

probability of the measurement accepting is the squared inner product of $|\phi_I\rangle$ and $|\psi\rangle$. This value is $(\sum_{i\in I:x_i=1} \alpha_i/n^{1/4})^2 \leq (1/\sqrt{n})(0.9\sqrt{n})(\sum |\alpha_i|^2) \leq 0.9$. Hence the protocol has soundness error 0.9, which can be improved by parallel repetition.